

REMARKS

On July 8, 2002, Applicants submitted a preliminary amendment including a substitute specification and proposed drawing corrections. Applicants request Examiner to confirm that the substitute specification and drawing corrections were entered.

The Examiner rejected claims 1, 17 and 28 under §112, second paragraph due to lack of antecedent basis for the limitation “bitfile”. Applicants amended claims 1, 17 and 28, along with dependent claims 2, 4, 14, 18, 20 and 29 to replace the term “bitfile data” with “file data,” which has proper antecedent basis in claims 1, 17 and 28.

The Examiner rejected claims 1-29 under 35 U.S.C. §102(e) as being anticipated by U.S. Pat. No. 6,282,548 of Burner (“Burner”). Applicants respectfully assert that the pending claims cannot be anticipated by Burner because they each include one or more elements that are not disclosed in Burner. For instance, in each of the claimed inventions, the metadata for each file contains information that correlates the location of that file to a file identifier. Claims 1, 16, 17 were amended to further clarify this point. When a file is requested, the metadata for that file is accessed and used to locate and retrieve the file. This is fundamentally different from Burner, where the file (e.g., web page) and metadata are located and retrieved and displayed *independently and concurrently*. (See e.g., Burner at col. 2, lines 48-59). This fundamental distinction is further described below.

Claims 1-15

Claim 1 includes elements that are not disclosed by Burner. Particularly, claim 1 recites a metadata service that maintains metadata corresponding to files, “a portion of the metadata for each file correlating a respective location of the file in the storage service with the respective file identifier.” Claim 1 further recites that the gateway service communicates with the metadata service “to determine the location of the file data in the storage service area in accordance with the file identifier.” Burner discloses neither of these elements.

In the system of claim 1 a file includes: (i) corresponding file data; (ii) a file identifier; and (iii) metadata correlating the location of that file (i.e., that file's file data) in storage with the file identifier. In order to access a file, the gateway service of the claimed system uses the metadata to locate the file based on its file identifier.

Burner includes a different type of metadata and operates in a fundamentally different manner. Burner teaches the use of metadata relating to the *contents* of the file (web pages), such as the author, the published date, related references, classifications, relevance ranking, etc. (Burner at col. 18, lines 20-65 and col. 19, lines 1-30). In fact, Burner generates the metadata by examining the *contents* of the files. (*Id.*) In Burner, the metadata for a file does *not* include information that correlates the location of that file to a file identifier, and the metadata is *not* used to locate and retrieve the contents of that file.

While the Examiner is correct that the metadata of a file or web page in Burner may include location information, the information does *not* provide information that is used to locate the page *itself*. Rather, the information relates to links to *other* pages. That is, in the cited section (col. 5, line 40-60, Burner describes keeping (possibly out-of date) pages that a web-crawler might have obtained. These pages are scanned to generate metadata based on the contents of the web-pages, including links to other web-pages that may be contained in this page.

The claimed system describes a different type of metadata, which is used in an entirely different way. In the present invention, the *contents* may be transparent to this system (hence the reference to the term "bitfile" or "file data"), and the metadata is data about how to access and fetch (i.e., retrieve) the *contents* or file data from the storage service (e.g., by the gateway service). In particular, the metadata in this system includes information about the location, in terms of which storage servers have the data, and information about the type of the storage used to save the contents. This metadata is generated when the contents (the file data) are saved, and is used when retrieving the file data at a later date. The metadata servers that serve the metadata, along with the storage servers that serve the file data, together constitute the symmetric

distributed file storage system. When a file is requested, the gateway server first communicates with the metadata service to obtain the location information for the contents of the file. Then the file contents or file data is retrieved using this metadata information.

Thus, important differences between Burner and the claimed invention are that Burner's system does not include metadata for a file that correlates an identifier for that file to the location of the actual contents of that file, and that Burner does not use a file's metadata to locate the contents of that file. In stark contrast, Burner teaches that a file (page) and its metadata are obtained *concurrently and independently*. As explained by Burner:

When a user, via the browser, asks to view a web page, the client software detects the request (or is notified of the request by the browser). **While the web browser is communicating with a web server to obtain the requested web page, the client communicates with the database server to obtain metadata about the requested page.** After the browser receives its requested information from the web server, it displays the requested web page in a conventional manner. **The client then displays its received metadata on the same display as the web page, and concurrently with the web page. Thus, the user can view the web page and metadata about the web page at the same time.**

(Col. 2, lines 48-59)(emphasis added)(*see also* col. 8, lines 14-25, which further describe this teaching). By teaching that a file and its metadata are concurrently and independently accessed and displayed, Burner directly contradicts and teaches away from the claimed invention, where the metadata is accessed first by the gateway server, which used the metadata to retrieve the file contents.

Thus, Burner does not teach either (i) metadata, “a portion of the metadata for each file correlating a respective location of the file in the storage service with the respective file identifier;” or (ii) a gateway service that communicates with the metadata service “to determine the location of the file data in the storage service area in accordance with the file identifier.” For at least these reasons, claim 1 and all claims depending from claim 1 (i.e., claims 2-15) cannot be anticipated by Burner.

Claim 14

Claim 14 is further patentable over Burner for the following independent reasons. The examiner rejects claim 14 as being anticipated by Burner, asserting that Burner discloses a hierarchical storage management scheme that allows the metadata service to cooperate with the storage service to manage a migration of file data from a first set of storage devices to a second set, as taught by claim 14. However, Burner's system does not discuss this type of storage for the files (web-pages) at all. In fact, Burner teaches that the *contents* are managed by a system external to the Burner system. As set forth in various locations throughout the specification (e.g., Burner col., 5, lines 50-55; col. 20, line 45, or at col. 15 lines 15-25), the Burner system copies the *contents* from the externally managed system into the Burner system for archival purposes, and the Burner system could be out-of-date or even contain pages that no longer exist in the external system.

The claimed invention provides different classes of storage for the same files. The storage devices are classified based on their capacity and performance characteristics. A file may move from one class of storage to another based on information about, say, its access pattern, without changing its file identifier. The metadata service “cooperates with the storage service to manage a migration” when a file is migrated to a different class of storage, “whereby clients requesting access to the moved files corresponding to the migrated file data need only supply the file identifiers corresponding to the moved files.” Thus, in the claimed invention, the metadata and actual file data are kept consistent at all times, so that the file identifier can always be used to locate the file.

In contrast, Burner merely describes an archiving process where *contents* are copied from an externally managed system into the Burner system, which could be out-of-date or even contain pages that no longer exist in the external system.

Because Burner does not teach a metadata service that manages a migration of file data that maintains consistency, allowing the file data to be accessed only by supplying the file identifier, it cannot anticipate claim 14 on this independent basis.

Claim 16

Like claim 1, claim 16 recites similar elements that are not disclosed in Burner. Particularly, claim 16 recites a metadata service that maintains metadata corresponding to files, “a portion of the metadata for each file correlating a respective location of the file in the storage service with the respective file identifier.” Claim 16 further recites that the gateway servers communicate with the metadata service “to determine the location of the bitfile data in the storage servers in accordance with the file identifier.”

As set forth above, Burner discloses neither of these elements. Therefore, for at least these reasons, claim 16 is patentable over Burner.

Claims 17-29

Like claim 1, claim 17 recites similar elements that are not disclosed in Burner. Particularly, claim 17 recites maintaining metadata corresponding to files, “a portion of the metadata for each file correlating a respective location of the file in the first storage device with the respective file identifier.” Claim 17 further recites “determining the location of the file data in the first storage device in accordance with the file identifier from the metadata maintained in the second device.”

As set forth above, Burner discloses neither of these elements. Therefore, for at least these reasons, claim 17 and all claims depending from claim 17 (i.e., claims 18-29) cannot be anticipated by Burner.

Claims 28-29

Claim 28 is further patentable over Burner for the following independent reasons. The Examiner rejects claim 28 as anticipated by Burner, citing col. 9, lines 40-45 of Burner. Those lines refer to the displaying of metadata on the same display device as the web-page. Claim 28 describes a technique to move portions of a file (some or all) from one storage device to another storage device, and this movement being transparent to the accessing client. The client continues

to use the file descriptor that it was provided earlier, without being aware that the storage for the file had actually changed. The metadata service is updated and becomes aware of the movement. In contrast, Burner explains that his metadata service can be well out of date, and might even provide data for non-existent web-pages. (See e.g., Burner, col. 5, lines 50-55; col. 20 line 45, or at col. 15, lines 15-25).

For all of these reasons and for the reasons set forth above relative to claim 14, Burner does not anticipate claim 28 or any claim depending from claim 28 (i.e., claim 29).

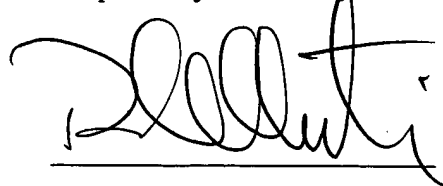
CONCLUSIONS

For all of these reasons, Applicants respectfully assert that all pending claims 1-29 are in condition for allowance. The Examiner's early reconsideration is respectfully requested. If the Examiner has any questions, the Examiner is invited to contact Applicant's attorney at the following address or telephone number:

David Alberti
c/o Patent Department
GRAY CARY WARE & FREIDENRICH LLP
2000 University Avenue
East Palo Alto, CA 94303-2248

Dated: October 22, 2004

Respectfully submitted,

A handwritten signature in black ink, appearing to read 'D. Alberti', written over a horizontal line.

David Alberti
Reg. No. 43,465